

CLAIMS

What is claimed is:

1. A persistent memory access system, comprising:
a primary region corresponding to a predefined portion of a primary network persistent memory unit (nPMU) communicatively coupled to at least one client processor node via a communication system, wherein the primary region is assigned to a client process running on the client processor node and is configured to store information received from the client process; and
a mirror region corresponding to a predefined portion of a mirror nPMU communicatively coupled to the client processor node via the communication system, wherein the mirror region is assigned to the client process and is configured to store the information received from the client process.
2. The system of claim 1, further comprising the primary nPMU and the mirror nPMU, wherein the primary nPMU and the mirror nPMU are physically separate units and are characterized by separate fault domains.
3. The system of claim 1, wherein the primary region comprises a plurality of primary virtual addresses corresponding a plurality of physical locations where the information is stored in the primary region, and wherein the mirror region comprises a plurality of mirror virtual addresses corresponding to another plurality of physical locations where the information is stored in the mirror region.
4. The system of claim 3, wherein the primary nPMU is configured to translate between one of the primary virtual addresses and a corresponding client address associated with the information, and wherein the mirror nPMU is configured to translate between one of the mirror virtual addresses and the corresponding client address associated with the information.

5. The system of claim 4, further comprising a base pointer, the base pointer corresponding to a difference in the primary virtual address and the corresponding client address such that the primary nPMU translates, and wherein the base pointer further corresponds to a difference in the mirror virtual address and the corresponding client address such that the mirror nPMU translates.

6. The system of claim 1, further comprising metadata residing in the primary nPMU and the mirror nPMU, wherein the metadata identifies the primary region and the mirror region assigned to the client process.

7. The system of claim 6, further comprising an attribute cache, the attribute cache residing in the client processor node and corresponding to at least some attribute information in the metadata.

8. The system of claim 1, further comprising a persistent memory unit (PMU) library residing in the client processor node, wherein the PMU library comprises at least a first function configured to permit an executing process to directly write the information to the primary region and the mirror region, and a second function configured to permit the executing process to directly read the information from a selected one of the primary region or the mirror region.

9. The system of claim 8, further comprising an application program interface (API) residing in the client processor node, wherein the API causes the client process to access the functions of the PMU library in response to a request by the client process to access the information.

10. The system of claim 1, further comprising a persistent memory manager (PMM) communicatively coupled to the client processor node via the communication system, wherein the PMM responds to a request from the client process for an assignment of persistent memory, and wherein the PMM causes the primary nPMU to create the primary region and causes the mirror nPMU to create the mirror region.

11. The system of claim 1, further comprising a persistent memory manager (PMM) communicatively coupled to the client processor node via the communication system, wherein the PMM responds to a request from the client process to end an assignment of persistent memory, and wherein the PMM causes the primary nPMU to delete the primary region and causes the mirror nPMU to delete the mirror region.

12. A client processor node comprising:
a processor;
a network interface communicatively coupled to a primary network persistent memory unit (nPMU), a mirror nPMU and a persistent memory manager (PMM) via a communication system; and

a memory, the memory further comprising:

a process that accesses a primary region corresponding to a predefined portion of the primary nPMU and that accesses a mirror region corresponding to a predefined portion of the mirror nPMU when the process is executed by the processor; and

a persistent memory unit (PMU) library, wherein the PMU library comprises at least a first function configured to permit the process to directly write information to the primary region and the mirror region, and a second function configured to permit the process to directly read the information from a selected one of the primary region or the mirror region.

13. The client processor node of claim 12, further comprising an application program interface (API) residing in the memory, wherein the API, when the API is executed by the processor, causes the client processor node to access the functions of the PMU library so that the process accesses the primary region and the mirror region.

14. The client processor node of claim 12, further comprising an attribute cache residing in the memory and corresponding to at least some access information in metadata, wherein the metadata identifies at least the primary region and the mirror region assigned to the client process.

15. A method directly accessing persistent memory in a primary network persistent memory unit (nPMU) and a mirror nPMU, the method comprising:

allocating a first region in the primary nPMU, the first region corresponding to a portion of memory in the primary nPMU;

allocating a second region in the mirror nPMU, the second region corresponding to a portion of memory in the mirror nPMU;

determining region information corresponding to the first region and the second region; and

storing the determined region information as metadata in a first metadata region in the primary nPMU and in a second metadata region in the mirror nPMU.

16. The method of claim 15, further comprising:

receiving an initial request from a client process for access to the persistent memory; and

creating the first region and the second region, before the allocating, determining and storing are performed.

17. The method of claim 15, further comprising:

receiving a subsequent request from a client process for access to the persistent memory; and

opening the first region and the second region.

18. The method of claim 17, further comprising communicating a base pointer value to the client process.

19. The method of claim 17, further comprising communicating a region size parameter to the client process.

20. The method of claim 17, further comprising:

receiving a write request communicated from the client process to the persistent memory system;

receiving information that is to be stored in the persistent memory system;

storing the information in the first region; and

storing the information in the second region.

21. The method of claim 17, further comprising:

receiving a read request communicated from the client process to the persistent memory system;

retrieving information that is stored in the persistent memory system based upon the region information specifying at least one virtual address corresponding to a first physical address where the information is stored in the first region and corresponding to a second physical address where the information is stored in the second region;

communicating the retrieved information from the persistent memory system to the client process.

22. The method of claim 21, further comprising:

receiving a mirror parameter in the read request, the mirror parameter specifying that the information is to be retrieved from the first region; and

retrieving the information from the first region.

23. The method of claim 22, further comprising:

determining if the information is available from the first region; and

retrieving the information from the second region when the information is not available from the first region.

24. The method of claim 21, further comprising:

receiving a mirror parameter in the read request, the mirror parameter specifying that the information is to be retrieved from the second region; and

retrieving the information from the second region.

25. The method of claim 24, further comprising:
determining if the information is available from the second region; and
retrieving the information from the first region when the information is not available from the second region.

26. The method of claim 17, further comprising:
receiving a request from the client process to end access to the persistent memory system; and
closing the first region and the second region.

27. The method of claim 15, further comprising:
receiving a subsequent request from a client process to permanently terminate access to the persistent memory system; and
deleting the first region and the second region.

28. The method of claim 15, further comprising:
receiving a request from a client process to receive access information corresponding to at least part of the metadata;
retrieving the requested access information from a selected one of the first metadata region in the primary nPMU and the second metadata region in the mirror nPMU; and
communicating the retrieved access information to the client process.

29. The method of claim 15, further comprising:
receiving a request for list information corresponding a plurality of regions in the primary nPMU and the mirror nPMU;
retrieving the requested list information from a selected one of the first metadata region in the primary nPMU and the second metadata region in the mirror nPMU; and
communicating the list information corresponding to the regions.

30. A method directly accessing a persistent memory system from a client processor node, the method comprising:

executing a client process that accesses the persistent memory system;

executing an application process interface (API), the API retrieving a region handle identifying a first region in a primary network persistent memory unit (nPMU) and identifying a second region in a mirror nPMU, retrieving at least one region function related to a function required by the executing client process, and retrieving at least one parameter associated with the function;

generating an access request comprising the region handle, the region function and the parameter; and

communicating the access request to the persistent memory system.

31. The method of claim 30, further comprising:

communicating an initial access request for access to the persistent memory system; and

receiving region information identifying a region, the region corresponding to the first region in the primary nPMU and corresponding to the second region in the mirror nPMU.

32. The method of claim 31, further comprising:

communicating a subsequent request for access to the persistent memory system;

specifying a region name; and

communicating information to be stored in the persistent memory system, wherein the location of the information is determined by the specified region and a base pointer value corresponding to a virtual address of the first region and the second region.

33. The method of claim 32, further comprising:

generating the region handle corresponding to the region information and the region name; and

identifying the region using the region handle in subsequent communications.

34. The method of claim 31, wherein the communicated access request comprises a base nPMU network virtual address and a network identification (ID).

35. The method of claim 31, wherein the communicated access request comprises a primary base nPMU network virtual address, a mirror base nPMU network virtual address, a primary network identification (ID) and a mirror network ID.

36. The method of claim 30, further comprising:
communicating a subsequent request for access to the persistent memory system;
specifying a region name corresponding to the first region and the second region,
a base pointer value corresponding to a virtual address of the first region and the second region, a read pointer parameter corresponding to the virtual address where reading information is to start, and a length parameter corresponding to an amount of the information to be read; and

receiving the information from the persistent memory system, wherein the location of the received information is determined by the specified region name, the base pointer value, the read pointer parameter and the length parameter.

37. The method of claim 30, further comprising:
communicating an attribute request to the persistent memory system;
receiving attribute information corresponding to a region; and
storing the received attribute information into an attribute cache.

38. A system for accessing a persistent memory system, comprising:

- means for receiving an initial request from a client process for access to the persistent memory system;
- means for allocating a first region in a primary network persistent memory unit (nPMU), the first region corresponding to a portion of memory in the primary nPMU;
- means for allocating a second region in a mirror nPMU, the second region corresponding to a portion of memory in the mirror nPMU;
- means for determining region information corresponding to the first region and the second region; and
- means for storing the determined region information in a first metadata region in the primary nPMU and in a second metadata region in the mirror nPMU.

39. The system of claim 38, further comprising:

- means for receiving a subsequent request from the client process for access to the persistent memory system; and
- means for opening the first region and the second region.

40. The system of claim 39, further comprising:

- means for receiving a write request communicated from the client process to the persistent memory system;
- means for receiving information that is to be stored in the persistent memory system;
- means for storing the information in the first region; and
- means for storing the information in the second region.

41. The system of claim 39, further comprising:

means for receiving a read request communicated from the client process to the persistent memory system;

access means for retrieving information that is stored in the persistent memory system based upon the region information in the read request specifying a region name corresponding to the first region and the second region, a base pointer value corresponding to a virtual address of the first region and the second region, a read pointer parameter corresponding to another virtual address where reading the information is to start, and a length parameter corresponding to an amount of information to be read;

means for communicating the retrieved information from the persistent memory system to the client process.

42. A system for accessing a persistent memory system by a client processor node, comprising:

means for executing a client process that accesses the persistent memory system;

means for executing an application process interface (API), the API retrieving a region handle identifying a first region in a primary network persistent memory unit (nPMU) and identifying a second region in a mirror nPMU assigned to the client processor node, retrieving at least one region function related to a function required by the executing client process, and retrieving at least one parameter associated with the function; and

means for communicating an access request to the persistent memory system.

43. The system of claim 42, further comprising:

means for communicating a request to the persistent memory system;

means for specifying a region name, the region name corresponding to the first region in the primary nPMU and corresponding to the second region in the mirror nPMU; and

means for communicating information to be stored in the persistent memory system, wherein the location of the information is determined by the specified region and a base pointer value corresponding to a virtual address of the first region and the second region.

44. The system of claim 42, further comprising:

means for communicating a request to the persistent memory system;

means for specifying a region name corresponding to the first region and the second region, a base pointer value corresponding to a virtual address of the first region and the second region, a read pointer parameter corresponding to the virtual address where reading information is to start, and a length parameter corresponding to an amount of information to be read; and

means for receiving the information from the persistent memory system, wherein the location of the received information is determined by the specified region name, the base pointer value, the read pointer parameter and the length parameter.

45. A computer-readable medium having a program for accessing a persistent memory system, the program comprising logic configured to perform:

executing a client process that accesses the persistent memory system;

accessing a persistent memory library, the persistent memory library comprising a region handle identifying a first region in a primary network persistent memory unit (nPMU) and identifying a second region in a mirror nPMU, at least one region function related to a function required by an executing process, and at least one parameter associated with the function;

executing an application process interface (API) that generates an access request comprising the API retrieving the region handle, the region function and the parameter; and

communicating the access request to the persistent memory system.

46. A computer-readable medium having a program for accessing a persistent memory system, the program comprising logic configured to perform:

- receiving an initial access request from a client process requesting access to the persistent memory system;

- allocating a first region in a primary network persistent memory unit (nPMU), the first region corresponding to a portion of memory in the primary nPMU;

- allocating a second region in a mirror nPMU, the second region corresponding to a portion of memory in the mirror nPMU;

- determining access information corresponding to the first region and the second region; and

- storing the determined access information in a first metadata region in the primary nPMU and in a second metadata region in the mirror nPMU.